

SPECIFICATION
(Sprint Docket No. 2334)

5

Title: **Implicit Floor Control in a Packet-Based Real-Time Media Session**

10

Inventors: Christopher M. Doran
Ryan H. Hodgson

15

Assignee: Sprint Spectrum L.P.
6391 Sprint Parkway
Overland Park, KS 66251-2100

20

BACKGROUND

1. Field of the Invention

The present invention relates to network communications and, more particularly, to the management of packet-based real-time media sessions.

5 2. Description of Related Art

As a general matter, it is known to establish a real-time media conference over a packet-switched network between multiple user stations, each operated by a respective user. A communication server, such as a multipoint conference unit (MCU) for instance, can reside functionally in the network and can operate as a bridging or switching device between the 10 participating stations, to support the conference session.

In practice, a participating station might initiate the conference session by sending to the communication server a session setup message that identifies the other desired participant(s). In response, the server may then seek to connect each of the designated other participants, such as by forwarding the session setup message or sending a new session setup message to each other 15 party. Ultimately, the server would thereby establish a conference leg with each participating station, including the initiating station, and the server would then bridge together the legs so that the users at the stations can confer with each other, exchanging voice, video and/or other media in real-time via the server.

A signaling mechanism such as the well known Session Initiation Protocol (SIP) could be 20 used to initialize the conference and more particularly to set up each conference leg. Further, digitized media could be packetized and carried between each participating station according to a mechanism such as the well known Real-time Transport Protocol (RTP), for instance. The core

industry standards for SIP (Internet Engineering Task Force (IETF) Request For Comments (RFC) 2543) and RTP (IETF RFC 1889) are hereby incorporated by reference.

Packet based media conferencing can be advantageously employed to provide an "instant connect" service, where a user of one station can readily initiate a real-time media conference 5 with one or more designated target users at other stations. The initiating user may simply select a target user or group and then press an instant connect button on his or her station, and the user's station would responsively signal to a communication server to initiate a conference between the initiating user and the selected user or group. This sort of service is referred to as "instant connect" because it strives to provide a quick connection between two or more users, in contrast 10 to telephone service where a user dials a telephone number of a party and waits for a circuit connection to be established with that party.

An example of an instant connect service is commonly known as "push-to-talk" (PTT). In a PTT system, some or all of the conference stations are likely to be wireless devices such as 15 cellular mobile stations, that are equipped to establish wireless packet-data connectivity and to engage in voice-over-packet (VoP) communication. Alternatively, some or all of the stations could be other sorts of devices, such as multimedia personal computers or Ethernet-telephones, that can establish packet data connectivity and engage in VoP communication through landline connections. Further, each station could be equipped with a PTT button or other mechanism that 20 a user can engage (actuate) in order to initiate an PTT session or to request the floor during an ongoing session.

In practice, a user of a PTT-equipped mobile station might select a target user or group of users from a contact list or other program menu and engage the PTT button to initiate a conference session with that user or group. In response, the mobile station may then send a

session initiation message to the communication server, to set up a conference session in the manner described above for instance, and the user could begin talking with the other users. Further, a similar mechanism could be applied to establish real-time media conferences carrying video or other media as well.

5 A conferencing system could be designed to provide either full-duplex service or half-duplex service. In a full-duplex system, a participating station would be allowed to send and receive media at the same time, so that a user of the station could both talk and listen at once. In order to accommodate full-duplex operation, a communication server would be configured to receive media from multiple stations at once and to output to each station a mixture of the media 10 or some representative subset of the media (e.g., a strongest signal).

15 In a half-duplex system, on the other hand, a participating station would at any time be allowed to either send media to the server or receive media from the server, but would be precluded from sending and receiving concurrently. In order to accommodate half-duplex operation, a communication server would be configured to apply a floor-control process, according to which the server allows only one station to have the floor at once. Thus, in a half-duplex mode, a participating station would receive media from only the handset that has the floor.

20 In a typical floor control process, a participant must request permission to "speak" (i.e., to send voice or other media) by sending a "floor-request" message to the server. The server then replies with a message that either grants or denies the floor. Once the server grants the floor to a participant, the server blocks all other participants from speaking (by denying all floor requests) until the speaker sends a "floor-relinquish" message to the server and the server acknowledges. Upon relinquishment of the floor, the server would then send a "floor-relinquished" message to

all participants and the participants would acknowledge. Only after this entire sequence has been completed will any other participant be allowed to speak.

Unfortunately, however, these floor control message exchanges can introduce delay into the communication process. Consequently, an improved floor control process is desired.

5

SUMMARY

The present invention provides an implicit floor control mechanism for packet-based real-time media communications. In accordance with an exemplary embodiment of the invention, when a user presses the floor-control button (e.g., a PTT button) on a participating station, the station will assume that it has the floor and begin transmitting media to the communication server. If the floor is currently open, then the server will grant the floor to the first station from which the server receives a media stream. And the server will ignore incoming media from any other participant while the speaker holds the floor.

In other words, a participating station will implicitly request the floor by simply beginning to send media to the server. If the floor is currently open, the server will treat the incoming media as an implicit floor request and will thus grant the floor to the requesting station. The server may then begin forwarding the speaker's media along to each other participating station.

By the same token, since communication is half duplex, a participating station can treat its receipt of media from the server as an implicit indication that it does not have the floor, i.e., as an indication that another participant has the floor. Therefore, if a station receives media from the server when the station is requesting the floor (e.g., by sending media to the server), the

station will treat the incoming media as an implicit denial of its floor-request. An exemplary station may then notify a user that the floor has been denied.

Advantageously, this implicit floor control process can thus reduce or altogether eliminate the exchange of express floor control requests and responses between participating stations and the communication server.

5

BRIEF DESCRIPTION OF THE DRAWINGS

An exemplary embodiment of the present invention is described herein with reference to the drawings, in which:

10 Figure 1 is a block diagram of a communication system in which the exemplary embodiment can be employed;

Figure 2 is a block diagram of an exemplary user station;

Figure 3 is a block diagram of an exemplary communication server;

Figure 4 is a block diagram of a wireless push-to-talk system;

15 Figure 5 is a message flow diagram depicting setup of a packet-based real-time media conference;

Figures 6-10 are flow charts depicting functions that an exemplary user station can carry out in accordance with the exemplary embodiment; and

20 Figure 11 is a flow chart depicting functions that an exemplary communication server can carry out in accordance with the exemplary embodiment.

DETAILED DESCRIPTION OF AN EXEMPLARY EMBODIMENT

1. Example Network Architecture

a. General

5 Referring to the drawings, Figure 1 illustrates a communication system 10 in which an exemplary embodiment of the present invention can be employed. For simplicity and by way of example, Figure 1 depicts three user stations 12, 14, 16 coupled with a common packet-switched network 18. User station 12 is operated by user A, user station 14 is operated by user B, and user station 16 is operated by user C. Sitting on the packet network 18, by way of example, are then a 10 proxy server 20 and a communication server 22.

It should be understood, of course, that this and other arrangements and processes described herein are set forth for purposes of example only, and other arrangements and elements (e.g., machines, interfaces, functions, orders of elements, etc.) can be added or used instead and some elements may be omitted altogether. Further, those skilled in the art will appreciate that 15 many of the elements described herein are functional entities that may be implemented as discrete components or in conjunction with other components, in any suitable combination and location, and by software, firmware and/or hardware.

In the arrangement of Figure 1, each user station 12, 14, 16 is preferably equipped with 20 hardware and logic to establish network connectivity, to set up and engage in packet-based real-time media conferences via server 22 and to engage in implicit floor control.

Figure 2 is a block diagram showing some of the components that an exemplary user station could contain in order to carry out these functions. As shown in Figure 2, the exemplary user station includes a processor 26, data storage 28, a user interface 30, and a communication interface 32, all of which may be coupled together by a system bus or other mechanism 34.

Each of these components may take various forms, the particular details of which are not necessarily critical. For instance, processor 26 may be one or more general purpose microprocessors (e.g., Intel Pentium class processors) or dedicated processors, either of which could integrate part or all of data storage 28. And data storage 28 may be volatile and/or non-volatile storage (such as RAM, flash memory and/or a storage drive).

5 User interface 30 may facilitate interaction with a user. As such, the user interface may include media input and output mechanisms. To facilitate voice communications, for instance, these mechanisms might include a microphone (not shown) for receiving analog speech signals from a user, and a speaker (not shown) for playing out analog speech signals to a user. (Further, 10 the mobile station will likely include digital/analog conversion circuitry (not shown) for converting between analog media signals and digital representations of those signals.)

In addition, the user interface 30 may include a display, speaker or other mechanism (not shown) for presenting information and menus to a user, as well as an input mechanism (e.g., keyboard, keypad, microphone, mouse, and/or touch-sensitive display overlay) (not shown) for 15 receiving input from a user. To facilitate floor control, the input mechanism may also include a floor-control button 36 or other mechanism that a user can readily engage in order to request the floor in an ongoing session.

Communication interface 32, in turn, facilitates communication through an access channel to packet network 18. The communication interface may thus vary in form depending 20 on the type of connection through which the station will communicate. For instance, if the station is coupled through a wired Ethernet connection to the network, then communication interface 34 might be a conventional Ethernet module. As another example, if the station is coupled through a wireless Ethernet or other radio access link to the network, then the

communication interface might include a suitable chipset and antenna for communicating according to a designated air interface protocol.

In the exemplary embodiment, data storage 28 may hold program logic, such as machine language instructions, that can be executed by processor 26 to carry out various functions 5 described herein. (Alternatively or additionally, the exemplary station could include hardware and/or firmware to carry out these functions.)

For example, to facilitate packet-data communications over network 18, the logic may define a conventional IP stack. As another example, to facilitate setting up and tearing down communication sessions, the logic may define a SIP user agent client application that enables 10 processor 26 to engage in conventional SIP messaging.

As still another example, to facilitate real-time media communication, the logic may define an RTP client application compliant with RFC 1889. And the logic may enable processor 26 to receive media signals from user interface 30 and to encode and packetize outgoing media as RTP/UDP/IP packets for transmission via communication interface 32 for receipt by server 15 22. Similarly, the logic may enable processor 26 to depacketize and decode incoming media signals provided by communication interface 32 from server 22 and to pass the decoded signals to user interface 32 for playout to a user.

In accordance with the exemplary embodiment, the logic may then further define mechanics for engaging in implicit floor-control as presently contemplated. In particular, the 20 logic may define mechanics for implicitly requesting the floor and for handling an implicit denial of the floor. For instance, when a user requests the floor, the logic may cause processor 26 to begin receiving media from the user and sending the media in an outgoing RTP stream to the server 22. And if the processor detects a user floor request at the same time as an incoming RTP

stream from server 22, the logic may cause the processor to treat the incoming RTP stream as a floor denial and to notify the user accordingly. Further details of this process will be described below.

Referring back to Figure 1, communication server 22 is preferably equipped with 5 hardware and logic to establish network connectivity, to set up and support packet-based real-time media conferences, and to engage in implicit floor control.

Figure 3 is a block diagram showing some of the components that an exemplary communication server 22 could contain in order to carry out these functions. As shown in Figure 10 3, the exemplary server 22 includes a processor 40, data storage 42, and a communication interface 44, all of which could be coupled together by a system bus or other mechanism 46.

As in the exemplary user station, each of these components may take various forms, the 15 particular details of which are not necessarily critical. For instance, processor 40 may be one or more general purpose microprocessors (e.g., Intel Pentium class processors) or dedicated processors, either of which could integrate part or all of data storage 28. And data storage 42 may be volatile and/or non-volatile storage (such as RAM, flash memory and/or a storage drive).

Communication interface 44 functions to provide connectivity with network 18. Like that in the exemplary user station, communication interface 44 may thus take various forms depending on the form of the link between the server 22 and network 18. By way of example, the communication interface 44 could be a wired or wireless Ethernet module.

20 Data storage 42, in turn, may hold program logic, such as machine language instructions, that can be executed by processor 40 to carry out various functions described herein. (Alternatively or additionally, the exemplary server 22 could include hardware and/or firmware to carry out these functions.)

Like the exemplary user station, for example, the logic may define a conventional IP stack to facilitate packet-data communications over network 18 and a SIP user agent client application to facilitate SIP messaging. The logic will also preferably define an RTP client application compliant with RFC 1889, as well as functionality to receive and forward RTP media streams.

Further, in the exemplary embodiment, the logic will further define mechanics for engaging in implicit floor-control as presently contemplated. In particular, the logic may define mechanics for granting the floor in response to an implicit floor request from a participating station, and for implicitly denying (i.e., ignoring) a floor request if another participant already has the floor. Further details of this process will be described below.

Additionally, data storage 42 would preferably hold a record of which, if any, station currently holds the floor at any moment. Thus, when server 22 grants the floor to a given station, processor 40 could record in data storage 42 that the station holds the floor.

b. Example Push-To-Talk Architecture

The arrangement shown in Figure 1 can generally represent any sort of communication system in which multiple stations can engage in packet-based real-time media communication with each other via a communication server. Regardless of how the participating endpoints gain network connectivity, the implicit floor control process of the exemplary embodiment can help streamline management of a half-duplex communication session.

By way of example, Figure 1 can represent a wireless PTT system, in which some or all of the user stations are wireless devices such as cellular mobile stations for instance. Figure 4 illustrates the arrangement of such a system 200.

As shown in Figure 4, the user stations are shown as cellular mobile stations 202, 204, 206. Each mobile station is served by a radio access network with an IP network 208. In particular, mobile station 202 is shown linked by a first radio access network 210 with the IP network, and mobile stations 204, 206 are shown linked by a second radio access network 212 with the IP network. Alternatively, each mobile station can be served by a separate radio access network, or all could be served by a common radio access network.

Each radio access network could take various forms, and the radio access networks may or may not be the same as each other. In the example shown in Figure 2, both radio access networks have the same configuration. In particular, each radio access network includes a base transceiver station (BTS) 214, 216 that can communicate with mobile stations over an air interface 219, 220, 221. Each BTS is then coupled with a base station controller (BSC) 222, 224, which may in turn be coupled with a mobile switching center (MSC) 226, 228 and with a packet data serving node (PDSN) 230, 232 or other gateway to the IP network 208. Other arrangements are possible as well.

Each mobile station may acquire radio connectivity and IP network connectivity in a manner well known in the art. For instance, applying well known "3G" recommendations, a mobile station may send an origination request over an air interface access channel to its MSC, and the MSC may forward the request back to the BSC. The BSC may then direct the mobile station to operate on a given traffic channel over the air interface. Further, the BSC may forward the request to the PDSN, and the PDSN may work with the mobile station to set up a data link, such as a point-to-point protocol (PPP) session between the mobile station and the PDSN. The PDSN may also assign a mobile-IP address to the mobile station, to allow the mobile station to engage in IP-network communications.

The air interface between each mobile station and its BTS preferably complies with an accepted protocol, examples of which include CDMA, TDMA, GSM and 802.11x. In the exemplary embodiment, for instance, the air interface protocol can be cdma2000, which is published by the 3rd Generation Partnership Project 2. Each mobile station may therefore be a 5 3G mobile station that is equipped to acquire wireless packet-data connectivity in a manner well known in the art.

Each mobile station is also preferably equipped to engage in SIP and RTP communication like the user stations described above. And each mobile station preferably includes a PTT button and associated logic, to allow a user to request the floor during an ongoing 10 session.

Further illustrated as nodes on IP network 208 are then a proxy server 234 and communication server 236, which are analogous to the proxy server 20 and communication server 22 in Figure 1.

2. Example Operation

15 a. Session Setup

The implicit floor-control process assumes that a packet-based real-time media session exists between two or more participating stations via a communication server. Figure 5 is a message flow diagram depicting an example of how such a session could be set up in the arrangement of Figure 1. It should be understood, however, that any other session setup process 20 could be used instead.

As shown in Figure 5, at step 50, user A directs station 12 to initiate a packet-based real-time media session (e.g., PTT session) with users B and C. For this purpose, the floor-control

button 36 on station 12 might also function as a session initiation button. So user A might engage that button.

At step 52, in response to the request from user A, station 12 sends a SIP INVITE to proxy server 20 for transmission in turn to server 22. The INVITE preferably designates users B 5 and C (or stations 14 and 16) or designates a group ID that server 22 can translate into users B and C (or stations 14 and 16).

At step 54, upon receipt of the INVITE, server 22 sends an INVITE to each target participant A and B, in an effort to set up an RTP conference leg with each target participant. At step 56, upon receipt of the INVITE from server 22, each target station signals its agreement to 10 participate, by sending a SIP "200 OK" message to server 22. At step 58, when the server 22 receives those messages, the server signals its agreement to participate by sending a 200 OK to station 12.

At step 60, station 12 then sends a SIP "ACK" message to server 22, to complete 15 signaling for setup of an RTP leg between station 12 and server 22. And at step 62, server 22 then sends an ACK to each target station to complete signaling for setup of an RTP leg between the server 22 and the target station. With the legs thus established, server 22 may then begin 20 bridging communications between the participating stations.

To begin with, station 12 may have the floor as a result of the fact that station 12 initiated the session. To facilitate a discussion of the implicit floor control process, assume that station 12 20 then relinquishes the floor, through express or implicit signaling with server 22. Thus, a packet-based real-time media session exists between stations 12, 14, 16 via server 22, and implicit floor control may proceed.

b. Implicit Floor Control at a User Station

As explained above, an exemplary user station will be arranged to carry out implicit floor control by sending implicit floor-requests and by detecting and responding to implicit floor denials. Figures 6-10 depict examples of these processes, which could be carried out by station 5 12, 14 or 16 in Figure 1.

Referring first to Figure 6, a basic implicit floor request process is shown. At step 70, the station receives a floor-request from a user. For instance, the user may engage in the floor-control button 36 or otherwise interact with user interface 30 to signal a request to begin "talking" (sending voice and/or other media), which the processor 26 would detect. For example, 10 processor 26 could receive an interrupt when the user presses the floor-control button 36. In response, at step 72, the station begins sending an RTP stream representing media to server 22, as an implicit floor request. For instance, processor 26 may begin receiving voice from the user, digitizing and encoding the voice, and sending a digital representation of the voice in an outgoing RTP stream to server 22.

15 If the exemplary station is receiving an incoming RTP stream from server 22 at the time the user requests the floor, the station may still carry out the basic process of Figure 5. Alternatively, in that scenario, the station may consider the incoming RTP stream an implicit denial of the user's floor request. This latter process is illustrated by Figure 7.

As shown in Figure 7, at step 74, the station receives a floor request from the user. In 20 response, at step 76, the station determines whether it is currently receiving an incoming RTP stream from the server 22. If it is not receiving an incoming RTP stream from the server, then, at step 78, the station begins sending an RTP stream representing media to server 22. On the other hand, if it is receiving an incoming RTP stream from the server, then, at step 80, it may treat the

incoming RTP stream as an implicit floor denial, i.e., as an indication that someone else already has the floor.

In the exemplary embodiment, when a station detects an implicit floor denial like this, the station can notify the user that the floor has been denied, and the station would preferably decline 5 to send media to the server until the floor is released. Figure 8 depicts an example of this process, where the station takes both of these actions in response to an implicit floor denial.

As shown in Figure 8, at step 82, the station first detects an implicit floor denial. In response, at step 84, the station alerts the user who requested the floor. For instance, the station 10 may notify the user through an audible, visual or vibratory alert via user interface 30. As a specific example, the station may beep and present a display screen message that indicates another participant has the floor. Also in response, at step 86, the station declines to send media to the server, i.e., declines to send an implicit floor request to the server.

On the other hand, it is also possible that the station could begin receiving an incoming RTP stream from the server 22 while the station is sending media to the server. This could occur 15 in a race scenario, for instance, where a station implicitly requests the floor before starting to receive media that the server just began receiving from another station. In that case, the station will preferably treat the incoming RTP stream as an implicit floor denial. Figure 9 depicts this process.

As shown in Figure 9, at step 88, the station receives a floor request from a user. At step 20 90, as in Figure 6, the station responsively begins sending an outgoing RTP stream to the server 22. At step 92, while the station is sending the outgoing RTP stream to the server, the station receives an incoming RTP stream to the server. In accordance with the exemplary embodiment,

at step 94, the station would treat the incoming RTP stream as an implicit denial of the floor request, i.e., as an indication that someone else already has the floor.

When the station detects an implicit floor denial like this, the station can then notify the user of the floor denial, and the station would preferably stop sending the outgoing RTP stream to the server. Figure 10 depicts an example of this process, where the station takes both of these actions in response to an implicit floor denial.

As shown in Figure 10, at step 96, the station detects an implicit floor denial. In response, at step 98, the station alerts the user who requested the floor. Also in response, at step 100, the station stops sending the outgoing RTP stream to the server.

Note that a station can also voluntarily release the floor when the user of the station is finished speaking, such as when the user releases a floor-control button or otherwise signals a desire to release the floor. To release the floor, the station could then responsively send a signal of some sort to the server (such as a predefined code or bit in an RTP header, for instance).

c. Implicit Floor Control at the Communication Server

As further noted above, the exemplary communication server 22 can be arranged to carry out implicit floor control by receiving implicit floor requests, granting the floor in response to an implicit floor request when the floor is currently open, and implicitly denying the floor by disregarding an implicit floor request when the floor is not currently open. Figure 11 depicts an example of this process.

As shown in Figure 11, at step 110, the exemplary server detects that it is receiving an incoming RTP stream from a user station (such as station 12, 14 or 16), which the server treats as an implicit floor request. In response, at step 112, the server determines whether the floor is

currently open. For instance, processor 40 may consult data storage 42 to determine whether any station currently holds the floor.

If the server determines that the floor is currently open, then the server responds to the implicit floor request by granting the floor to the requesting station (i.e., station and/or user), and 5 updating its records in data storage 42. Thus, at step 114, the server may update its records in data storage 42 to reflect that the requesting station has the floor, and, at step 116, the server may begin to forward the media carried by the incoming RTP stream to each other participating station in an outgoing RTP stream.

On the other hand, if the server determines that the floor is not currently open, i.e., that 10 another station currently holds the floor, then, at step 118, the server implicitly denies the floor request by disregarding the incoming RTP stream (i.e., not forwarding the media and not responding to the requesting station). In other words, the server does not forward the media in the incoming RTP stream to the other participant(s). Optionally, the server may also send an express floor-denial message to the requesting station.

15 Note also that when the floor is released, the server can send a signal of some sort to each participating station, to alert the station(s) that the floor has been released. (For instance, the signal could be a predefined code or bit in an RTP header.)

3. Operation in a Full-Duplex Session

As described above, the exemplary embodiment is particularly advantageous in a half-20 duplex conference session, where a communication server outputs media at any given time from only the station that has the floor. However, the exemplary embodiment can be extended to apply in a full duplex session as well, provided that the notion of a "floor" exists in the session.

In a full-duplex session, as noted above, a communication server may output media from more than one participating station concurrently, so that participants can effectively talk and listen at the same time. In particular, the server could receive media from multiple participating stations, mix the underlying media together to produce combined media, and send to the 5 participating stations a media stream that embodies the combined media. Participating stations would then be equipped to handle the simultaneous input and output of media.

In a full-duplex session, the communication server could still be arranged to grant the floor to just one participant at a time. Having the floor in the full-duplex session, however, could have different meaning than having the floor in a half-duplex session. For instance, rather than 10 outputting media from only the station that has the floor (as in a half-duplex session), the communication server could be arranged to output media more loudly from the station that has the floor than from each other participating station. That is, when the server mixes together the media from various participating stations, the server could attenuate the media from each station that does not hold the floor or could amplify the media from the station that holds the floor.

15 Thus, applying the exemplary embodiment in the full-duplex scenario, the server may still grant the floor to a station in response to receipt of a media stream from the station, provided that no other participating station currently holds the floor. And the server may refuse to grant the floor to a station in response to receipt of media from the station, if another participating station currently holds the floor.

20 Further, the server could even be arranged to grant levels of floor to various participants in a full-duplex session. For example, the server could output most loudly the media from a station with a highest floor level, and the server could incrementally attenuate the media that it outputs from each other participating station having a successively lower floor level.

According to the exemplary embodiment, the server can accomplish this by granting the highest floor level to the first station from which the server receives a media stream when no other station currently holds the floor, granting the next floor level to the next station from which the server receives a media stream when the first station currently holds the highest floor level, 5 and so forth. As a station releases its control over a given floor level, the server could then responsively increment the floor levels of other stations that hold some level of the floor.

4. Conclusion

An exemplary embodiment of the present invention has been described above. Those skilled in the art will understand, however, that changes and modifications may be made to this 10 embodiment without departing from the true scope and spirit of the present invention, which is defined by the claims.